arXiv:1510.00010v1 [quant-ph] 30 Sep 2015

When is simpler thermodynamically better?

Andrew J. P. Garner,[1,2,*] Jayne Thompson,[1] Vlatko Vedral,[3,1,2] and Mile Gu[2,1,†]

[1] *Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, 117543, Singapore*
[2] *Center for Quantum Information, Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, 100084, China*
[3] *Atomic and Laser Physics, University of Oxford, Clarendon Laboratory, Parks Road, Oxford, OX1 3PU, United Kingdom.*

(Dated: October 2, 2015)

Living organisms capitalize on their ability to predict their environment to maximize their available free energy, and invest this energy in turn to create new complex structures. Is there a preferred method by which this manipulation of structure should be done? Our intuition is "simpler is better," but this is only a guiding principal. Here, we substantiate this claim through thermodynamic reasoning. We present a new framework for the manipulation of patterns—structured sequences of data—by predictive devices. We identify the dissipative costs and how they can be minimized by the choice of memory in these predictive devices. For pattern generation, we see that simpler is indeed better. However, contrary to intuition, when it comes to extracting work from a pattern, any device capable of making statistically accurate predictions can recover all available energy.

Simpler is better—an idea that has resonated with culture throughout history. Amongst many examples, the value of simplicity has been recognized in the minimalist movement in 1960s art and music, *shibui* in Japanese ceramics, productivism in post-revolutionary Russian, and the guiding principles of twenty-first century Apple product design [1]. Our predilection for simplicity even permeates natural philosophy, as most-famously formalized by William of Ockham's razor. This notion, "Everything should be made as simple as possible but no simpler," [2] suggests that satisfying explanations for reality must avoid unnecessary complications.

The virtue of simplicity in art is subjective, but in science this can be quantified, provided that we explain what is meant both by "simpler" and by "better". In thermodynamics, a subject whose original purpose relates to the design of steam-engines, "better" suggests efficiency and thrift (e.g. more power for less coal). This carries through to modern thermodynamics: the best approach for a given task minimizes the expenditure of a limited resource [3] (e.g. work or tricky-to-create states).

Meanwhile in computational mechanics [4, 5], we can formalize "simpler" in the context of pattern generation and manipulation. Everything we observe in the environment can be considered to be pattern—a temporal sequence of data that exhibits statistical structure. Much of science then deals with the construction of machines that take information from past observations, and use it to generate future statistics. Here, the *simplest* machine stores this information about the past using the minimal amount of internal structure.

To answer "When is simpler thermodynamically better?" we present a thermodynamic framework for pattern manipulation. We evaluate the steady-state operation of a full thermal cycle, consisting of a *generator* that invests work to accurately produce a pattern and an *extractor* that anticipates the pattern to exploit it for free energy (see fig. 1). We examine the work exchange in these devices, and how they use internal memory to retain their predictive ability. We then identify dissipative work costs that worsen their thermodynamic performance. In this context, a *simpler* device uses less internal memory, and a *better* one wastes less work maintaining this memory.

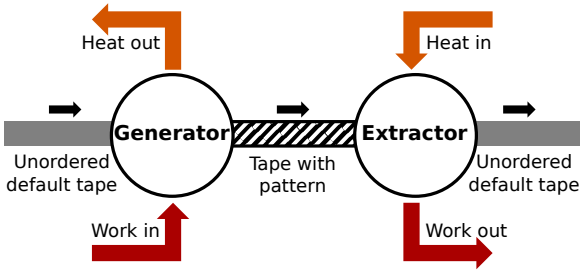

FIG. 1: **Cycle of pattern generation and extraction.** *A tape moves from left to right. A generator expends work to write a pattern to the tape. The extractor then uses the pattern on this tape to extract work. To run cyclically, each device maintains* prescient *memory that keeps track of the pattern. In this article, we identify the dissipative work costs. We find that simpler is thermodynamically better* only *when it comes to generating patterns.*

Ockham's razor is a guiding philosophical principle, rather than a cast-in-stone physical law; and so there is no a priori physical motivation to assume that using the smallest, simplest, memory is always thermodynamically better. By minimizing the dissipative work costs, we identify the best choice of internal memory. In the generator, using the simplest possible internal memory (the causal states associated with the pattern's *statistical complexity* [4, 5]) minimizes dissipation. For generating a pattern: simpler is thermodynamically better. For the extractor, defying our intuition, it transpires that the choice of memory has no thermodynamic consequence.

*Electronic address: ajpgarner@nus.edu.sg
†Electronic address: ceptryn@gmail.com

Patterns are a resource. Knowing a system's internal state has thermodynamic consequence: this knowledge can be used to perform work (drive a mechanical task), as illustrated by the Szilárd engine thought experiment [6]. A box has a single particle inside, on the left- or right-hand-side. A movable barrier inserted in the center acts as a piston that expands as the particle pushes against it. If an agent knows which side of the barrier the particle is on, she can couple the barrier (e.g. via a pulley) to raise a weight. As the piston expands it generates an amount of work $k_{\mathrm{B}}T\ln 2$ (at temperature T) by drawing in the same amount of heat from its surroundings.

Knowledge about patterns may also be exploited. Suppose an agent attempts to extract work from a series of Szilárd engines (labelled sequentially by index t), prepared such that the particle in engine t is on the same side of the barrier as in engine t–1 with probability $p \neq \frac{1}{2}$. An agent unware of this pattern would only be able to correctly predict the particle's location half of the time, and hence will extract less work than an agent who knows the pattern and couples her pulley accordingly. As such, the ability to predict grants thermodynamic advantage.

This is a manifestation of Maxwell's dæmon—an apparently paradoxical conversion of heat into work that is only resolved by accepting that information is physical and hence subject to the laws of thermodynamics [6–9]. For the single Szilárd engine, we must also account for the cost of resetting the agent's memory about the particle's location—this knowledge must be thought of as a resource. Likewise, since it is more thermodynamically useful that the sequence follows a pattern than be uncorrelated, the pattern itself must also be considered as a resource. Producing a pattern hence requires an investment of work. Moreover, any physical device that generates (or exploits) a pattern contains some memory about what has happened in the pattern so far, in order to accurately generate (or anticipate) upcoming parts of the pattern. Any thermodynamic costs of maintaining this internal memory must also be accounted for.

The quantitative link between information, entropy and heat dissipation is given by Landauer's principle [6]: the minimum work cost of any information-processing task is proportional to the change in information entropy[1] (just like macroscopic thermodynamics, where the work required to slowly change between two states of the same internal energy is proportional to the change in thermodynamic entropy).

The mathematics of patterns. We approach this mathematically using the language of statistical complexity [4, 5]. The outputs of a physical process are represented using a bi-infinite sequence of random variables $\ldots X^{t-1}X^tX^{t+1}\ldots$, where X^t describes the random variable at time t. (In our example above, X^t is whether engine t has a particle on the left or the right). These are partitioned such that variables up to X^t are in the past $\overleftarrow{X}^t$, and from X^{t+1} onwards are in the future $\vec{X}^t$. A particular *sequence* is specified by the values that each random variable ultimately takes ($X^1 = x_A$, $X^2 = x_B$ etc.). A statistical description of the *pattern* is given by the probability distribution over these sequences $\mathrm{P}(\overleftarrow{X}^t, \vec{X}^t)$.

When the outputs follow a pattern, the past and future are correlated. If a particular sequence $\overleftarrow{x}$ is observed in the past, the statistics of the future are given by $\mathrm{P}(\vec{X}^t \,|\, \overleftarrow{X}^t = \overleftarrow{x})$. Information theory quantifies how useful this past knowledge is for predicting the future: the difference in the entropy with and without knowledge of the past, $H(\vec{X}^t) - H(\vec{X}^t \,|\, \overleftarrow{X}^t)$. This value is exactly the *mutual information* between the past and future $I(\overleftarrow{X}^t\,; \vec{X}^t)$ (also known as the *excess entropy* E).

We make the simplifying assumption that the pattern is a *stationary process* [5]– the statistics $\mathrm{P}(\overleftarrow{X}^t, \vec{X}^t)$ are invariant under time translation. This does not mean that every output x^t in the sequence is identical, or that the pattern is Markovian, but rather that the statistics of X^{t+1} onwards, given a past sequence, have no explicit time dependence: $\mathrm{P}(\vec{X}^t \,|\, \overleftarrow{X}^t = \overleftarrow{x}) = \mathrm{P}(\vec{X}^{t'} \,|\, \overleftarrow{X}^{t'} = \overleftarrow{x})$ for all t, t' and $\overleftarrow{x}$. Hence, we can often omit the superscript t.

The notion of an agent is formalised by constructing a *prescient device*: a machine whose knowledge of future statistics is as good as is causally possible (i.e. without knowledge about the outcomes of random events in the future). It would be cumbersome to store the entire infinite past string in this device's memory, since this requires an unbounded amount of storage. Instead, one uses a (many-to-one) function to map multiple past sequences to the device's internal memory R, which will itself be a random variable over a finite alphabet. Two different pasts are only mapped to the same internal memory state if the future statistics conditional on them are the same. For the purpose of predicting future statistics, knowing R is as useful as knowing the entire past and $\mathrm{P}(\vec{X} \,|\, R) = \mathrm{P}(\vec{X} \,|\, \overleftarrow{X})$. Such internal states R are known as *prescient states*.

Storing these states requires an amount of memory given by the entropy $H(R)$ over their probability distribution. We define states that use less memory, and hence have lower entropy, as *simpler*. The simplest prescient states requiring the least memory are known as *causal states* (written S). These correspond to the equivalence classes of the pasts that predict the same future statistics. As any given pattern has a unique set of causal states, the memory $H(S)$ needed to store these states quantifies how difficult it is to describe that pattern, and so is a *statistical measure of complexity* [5]. Unlike Kolmogorov complexity [10] which concerns the generation of a specific string, the statistical complexity avoids classifying structureless but highly-random processes as complicated. Statistical complexity has been applied to analysing structure in diverse contexts, such as neural networks [11], financial markets [12] and cyber-security [13].

[1] Classically, given by the *Shannon entropy* $H(X) = -\sum_i P(X = x_i)\log_2 P(X = x_i)$. Using a base 2 logarithm gives units of *bits*.

The intuition bestowed upon us by Ockham's razor suggests that in any realization of a predictive device, we should use these simplest causal states. In the remainder of this article, we shall see where this intuition holds, and substantiate the penalties incurred by deviation from it.

A new framework for pattern manipulation. The classical expressions of thermodynamic laws (e.g. Kelvin's statement that a device cannot convert heat into work with no other effect [14]) concern cyclic behaviour—processes that leave the system in a state allowing for repetition with the same thermodynamic consequence. Without a full cycle in mind, there is the danger that the thermal benefit of a process may come at the expense of consuming an unaccounted-for resource. A cycle does not require the *microstate* of the system to return to its original value. Consider a piston of gas expanding and compressing: it does not matter if the individual molecules have moved to new locations by the end of the cycle, as long as the important thermodynamic variables—the pressure and volume—return to their original values.

In the same way, we present a cycle of generating and consuming a pattern (see fig. 1). Visualise the pattern as a series of symbols on a tape (each position labelled by t). These symbols are initially distributed according to uncorrelated high-entropy[2] random variables $X_{\text{dflt}} \ldots X_{\text{dflt}}$. The tape is acted on by two machines, which are in contact with a thermal reservoir (i.e. heat bath at inverse temperature $\beta = \frac{1}{k_{\text{B}}T}$) and a battery for storing free energy (e.g. a raising weight). The first machine, a *generator*, writes k steps of the pattern onto the tape with an associated work cost. The second machine, an *extractor*, resets k symbols on the patterned tape back to the uncorrelated default states, outputting work as it does so. An *instantaneous* device with $k = 1$ acts on one symbol at a time; whereas a machine with a larger *stride* k processes a larger section of the pattern at once.

Since the generator and extractor require accurate statistical knowledge of the future (the generator to know what it must produce next; the extractor to anticipate the upcoming inputs), they must both maintain prescient internal memory. As such they are realisations of prescient devices. We shall assume that these devices begin with their internal memories synchronized, such that the part of the pattern the generator is about to produce matches with the part of the pattern the extractor anticipates acting upon [i.e. the machines have internal states $R^t = r$ and $\tilde{R}^t = r'$ respectively such that $\mathrm{P}(\vec{X} \,|\, R = r) = \mathrm{P}(\vec{X} \,|\, \tilde{R} = r')$.] To run continually in an cycle, the generator and extractor must remain synchronized after they have generated and extracted the same number of symbols. This ensures that their performance on the next part of the pattern remains the same.

[2] The mathematical results presented will hold regardless of the particular statistics used for X_{dflt}, so long as we use the same statistical states for the generator's input and extractor's output.

Thus, after each device has processed k symbols, the prescient state must be updated from R^t to R^{t+k}. How this update is performed can be deterministic or otherwise. The behaviour is *deterministic* if given an initial memory state R^t and the next k symbols of the pattern $X^{t+1} \ldots X^{t+k}$, the final state R^{t+k} is completely defined (i.e. the randomness is entirely confined to the outputs $X^{t+1} \ldots X^{t+k}$). The alternative is that there is some additional randomness in the final value of R^{t+k}. However, this randomness can only be between choices of R^{t+k} that predict exactly the same future statistics of X, as otherwise the devices will cease to be prescient (and may desynchronize).

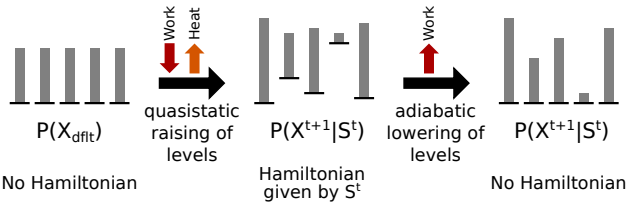

FIG. 2: **Writing a pattern to a tape.** *The various choices of symbols on a tape can each be associated with a different energy level of a system (drawn as a black horizontal lines whose relative height indicates relative energy). The statistical state of the symbol is a probability distribution (drawn as grey bars) over these configurations. By changing the Hamiltonian of the tape whilst remaining in contact with a thermal reservoir, the statistics can be altered to* $\mathrm{P}(X^{t+1} \,|\, S^t)$*. At this point, the system is isolated from the heat reservoir and the Hamiltonian is adiabatically removed. The whole procedure requires an investment of work proportional to the reduction in the state's entropy.*

Investing work to generate a pattern. For a given pattern, there is a family of generators, characterized by the prescient states R^t used as internal memory, and by the the number of steps k of the pattern that are generated at once. The work cost of generating a pattern consists of two contributions: one from changing the entropy of the tape, and one from updating the memory.

From Landauer's principle, the minimal work investment required to write k symbols of the pattern onto the tape (at inverse temperature $\beta = \frac{1}{k_{\text{B}}T}$) is given by

$$\beta\, W^k_{\text{tape}} = k \left[H(X_{\text{dflt}}) - H(X^{t+1} \,|\, S^t) \right], \qquad (1)$$

[See fig. 2, and *Appendix*.] This value relates to change in entropy of the tape itself, and so is an intrinsic property of the pattern rather than of the machine generating it. As such it can be expressed in terms of the causal states S (which are unique for any given pattern [5]) rather than the arbitrary device-dependent internal states R.

The second contribution, the cost of updating the internal memory into R^{t+k} (so that the generator is ready to produce the next part of the pattern) is given by

$$\begin{aligned} \beta W^k_{\text{diss}} = \quad & H(R^t \,|\, X^{t+1} \ldots X^{t+k} R^{t+k}) \\ & - H(R^{t+k} \,|\, R^t X^{t+1} \ldots X^{t+k}). \end{aligned} \qquad (2)$$

[See fig. 3 and *Appendix.*] The first term is the cost of erasing the previous state R^t, offset by the mutual information that is contained within the new state of the memory R^{t+k} and the patterned outputs $X^{t+1}\ldots X^{t+k}$. The second term (which is zero for deterministic updates) reflects the fact that if the update is indeterministic, according to Landauer's principle we can recover a portion of the work cost associated with the memory's change in entropy. However, by expanding the first term (details in *Appendix*), we arrive at

$$\begin{aligned}\beta W^k_{\text{diss}} = \quad & H(X^{t+1}\ldots X^{t+k}\,|\,R^t) \\ & - H(X^{t+1}\ldots X^{t+k}\,|\,R^{t+k}),\end{aligned} \tag{3}$$

where the indeterministic term has been eliminated. Whatever work might have been gained by introducing randomness into R^{t+k} is entirely cancelled out by the cost of resetting this randomness in the previous state R^t. Hence determinism does not in fact play a role in deciding the thermodynamic advantage of prescient memory.

Rather, we see that the dissipative cost is proportional to the difference between R^t's predictive power to guess the next k symbols, and R^{t+k}'s *retrodictive* power to guess the preceding k symbols. Failure to predict increases the first entropy, and hence the amount of dissipation. On the other hand, failure to retrodict increases the second entropy and lessens the total dissipation.

What choice of prescient memory R minimizes this cost? The first term $H(X^{t+1}\ldots X^{t+k}\,|\,R^t)$ is the same for *all* choices of prescient states, since they all must reduce the entropy of the future statistics by the largest possible amount. The second term $H(X^{t+1}\ldots X^{t+k}\,|\,R^{t+k})$ takes its largest value when R^{t+k} contains the least information about the preceding k symbols. The causal state S^{t+k} contain the least information about the past

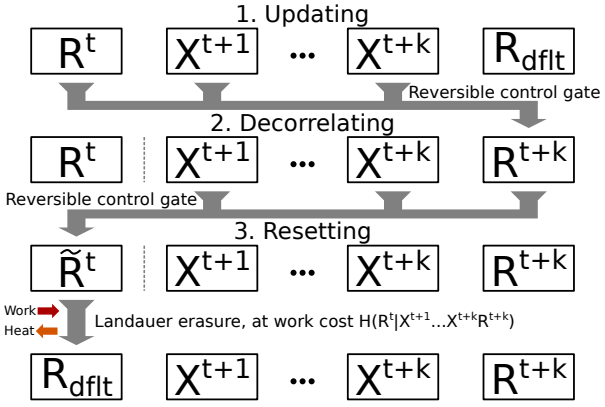

FIG. 3: **Updating the generator's memory.** *(The deterministic case is shown.) A blank ancilla state* R_{dflt} *is* updated *at no cost to* R^{t+k}*, conditioned on the values of the initial internal state* R^t *and outputs* $X^{t+1}\ldots X^{t+k}$*. The old state* R^t *is then* decorrelated *from* R^{t+1} *and* $X^{t+1}\ldots X^{t+k}$*, and the mutual information used to reduce the entropy of the internal state (transforming it into* $\tilde{R}$*). Finally* $\tilde{R}$ *is* reset *back to the blank ancilla state at work cost* $H(R^t\,|\,X^{t+1}\ldots X^{t+k}R^{t+k})$*, so that the generator's internal state is ready to produce the next part of the pattern.*

required to predict the future, so all alternative prescient states must contain this amount or more. As such, $H(X^{t+1}\ldots X^{t+k}\,|\,R^{t+k}) \leq H(X^{t+1}\ldots X^{t+k}\,|\,S^{t+k})$. Out of all the possible prescient states R to use as our internal memory, the simplest ones—the causal states S—hence dissipate the least heat. We thus conclude that for generating a pattern, *simpler is thermodynamically better.*

Extracting work from a pattern. Let us now evaluate how much work we can extract from a pattern, by considering the prescient *extractor*, This device is a sophisticated Szilárd engine [6] that maintains an up-to-date internal memory (the prescient state R^t) allowing it to anticipate the upcoming k symbols in the pattern. Work is extracted by converting $X^{t+1}\ldots X^{t+k}$ into uncorrelated output states $X_{\text{dflt}}\ldots X_{\text{dflt}}$.

In order to fully account for all the changes in entropy in our system, again we must consider both the tape on which the pattern is written and the internal memory of the extractor. We do this in a single step; the total average work released by taking k symbols to the default state $X_{\text{dflt}}\ldots X_{\text{dflt}}$, and updating the internal memory from R^t to $R^{t+k}$is given by

$$\beta W^k_{\text{out}} = k\left[H(X_{\text{dflt}}) - H(X^{t+1}\,|\,S^t)\right] \tag{4}$$

[Proof in *Appendix.*]

The value here is entirely proportional to the change in entropy of the k symbols on the tape, and is the maximum amount we would expect from Landauer's principle. The curious result here is that although we have accounted for the cost of updating, there is absolutely no dependence on the choice of internal memory R^t, and the scaling with the number of symbols processed k is trivial. Unlike with the generator, it does not matter what sort of memory is used for extraction.

Moreover, we remark that there may always be dissipation in generation since even causal states can contain superfluous information about the past for some patterns $\left[\text{in general } H(S) \geq I(\overleftarrow{X}\,;\overrightarrow{X})\right]$. However, this is not the case for extraction, even though we would expect larger R^t to perform worse since there should be more about the past to "clean up" when updating the internal state.

This apparent paradox may be explained by carefully considering the difference between the generator and the extractor. Crucially, the output pattern must be undisturbed at the end of the generation procedure but not for extraction. Thus, the difference is between a device that moves information, and a device that copies information. Moving information is a logically reversible process, whereas copying information is not—and it is *logical irreversibility* that lies behind dissipative costs in computation [7]. At the end of the extraction, there is only one copy of this information (encoded in R^{t+k}), whereas in generation, the information is present both in the internal memory and in the tape. Thus the extractor can move the information from $X^{t+1}\ldots X^{t+k}$ into R^{t+k}, but the generator must copy it. If more information is stored in R^t, then more has to be copied. This subtle, but important, distinction reveals to us why updating the memory

must dissipate heat in pattern generation, but not for pattern extraction.

Relation to existing results. We first observe that if we consider only updating the generator's memory, and restrict ourselves to causal states and the limit $k \to \infty$, then the dissipative work cost of updating the generator's memory [eq. (3)] converges on $I(\overleftarrow{X}; \vec{X}) - H(S)$, recovering the result of Wiesner et al. [15].

Secondly, we note that the dissipative cost [eq. (3)] may be re-expressed in terms of mutual information:

$$\beta W^k_{\text{diss}} = I(X^{t+1}\ldots X^{t+k}; R^{t+k}) - I(X^{t+1}\ldots X^{t+k}; R^t). \tag{5}$$

In the opposite, instantaneous, case $k = 1$, this recovers a mathematical quantity similar to that introduced by Still et al. [16] as the *useless instantaneous nostalgia*. That this coincides with our result in certain limits, despite very different derivations, demonstrates the universality of the information-theoretic concepts underlying the thermal costs of prediction. In both cases, the need to clean up unnecessary information about a process's past has been identified as a source of dissipation.

Our work extends this by examining the thermodynamics of a full cycle of pattern generation and extraction (see fig. 1). By doing so, we have identified that penalty for storing more than is necessary depends very much on how the prediction is used. Any prescient extractor can recover the entire free energy of the pattern—such a device has no concept of useless nostalgia, and no dissipation. It is not unnecessary storage in and of itself that induces the dissipative cost, but rather logical irreversibility. Only when the pattern manipulation is irreversible (such as in our generator, or the system in [16]) does storing too much information incur a thermal penalty.

The thermodynamics of patterns. The two devices we have presented in this article are building blocks that can be combined into complex pattern manipulators. The simplest example is the cycle of generation followed by extraction at the same temperature (fig. 1). This configuration could be considered as charging a battery that stores energy in the form of a pattern, to be later released by the extractor. Since the contributions from writing and consuming the pattern cancel out, the net cost is from the work dissipation whilst updating the generator's memory. Using the simplest internal memory ensures that the least work is wasted.

An alternative configuration is for a generator and extractor to act in parallel on different patterns. Here, the extractor consumes one pattern, and uses the energy released to power a generator that writes a different pattern. This approximates the action of all living organisms: for example, a lion metabolizes the structure of an antelope (destroying it in the process), and uses the energy released to build more lion. Using the simplest internal memory for generation grants the advantage that less antelope needs to be consumed in order to produce the same amount of lion.

Outlook. In this article, we have treated the thermodynamics of patterns, and identified the dissipative work costs in a cycle of producing and consuming a pattern. For generation, simpler is thermodynamically better; the cost is minimized when using the simplest internal states possible: the causal states associated with the pattern's statistical complexity. On the other hand, for extraction we found that any prescient device has the ability to recover the entire free energy of the pattern, because the internal memory could be updated in a logically reversible way. The answer to "When is simpler thermodynamically better?" is that simpler is better when it helps us avoid logical irreversibility.

Our discussion has thus far been restricted to the average behaviour of classical systems operating in the quasistatic regime. By taking the quasistatic limit, we do not account for additional dissipation that occurs when writing to the tape in finite time. The additional costs of this are highly implementation-dependent: typically, faster processes dissipate more heat. By considering only the average behaviour, we have avoided addressing the fluctuations in the work costs. Fluctuations are intrinsically related to dissipation and dissipative work [9, 17, 18], but it is not yet clear what the associated fluctuation relation is for the dissipative work discussed here.

Recent research indicates that by using quantum states, less memory is required than the best classical alternative for the same predictive power [19]. In a random process, there is a limit to how distinct the futures may be: two different pasts may in some (but not all) cases lead to exactly the same future. There is hence no need to perfectly distinguish between the two internal states, but this is unavoidable when using classical memory. Quantum memory does not require this, and hence has a lower entropy for the same level of prescience as the classical case. It is therefore feasible that by storing less unnecessary information about the past, quantum machines could be simpler devices and hence dissipate less heat. This would suggest that there are certain pattern-processing tasks where the quantumness of a device could yield thermodynamic advantage. Could quantum pattern generators, through their increased simplicity, be better than the classical alternatives?

Acknowledgements. We are grateful for funding from the John Templeton Foundation Grant 53914 *"Occam's Quantum Mechanical Razor: Can Quantum theory admit the Simplest Understanding of Reality?"*; the National Research Foundation; the Ministry of Education in Singapore, the Academic Research Fund Tier 3 MOE2012-T3-1-009; the National Basic Research Program of China Grants 2011CBA00300 and 2011CBA00302; the National Natural Science Foundation of China Grants 11450110058, 61033001 and 61361136003; the 1000 Talents Program of China; and the Oxford Martin School. We thank Blake Pollard, Merlin Seller, and Whei-Yeap Suen for illuminating discussion.

[1] S. Richmond. Jonathan Ive interview: simplicity isn't simple. *The Daily Telegraph (UK)*, May 2012. ISSN 0307-1235. http://www.telegraph.co.uk/technology/apple/9283706/Jonathan-Ive-interview-simplicity-isnt-simple.html.

[2] Attributed to A. Einstein via R. Sessions. Everything should be made as simple as possible, but not simpler, 2011. http://quoteinvestigator.com/2011/05/13/einstein-simple/.

[3] F. G. S. L. Brandão, M. Horodecki, J. Oppenheim, J. M. Renes, and R. W. Spekkens. Resource Theory of Quantum States Out of Thermal Equilibrium. *Physical Review Letters*, 111(25):250404, December 2013. ISSN 0031-9007. doi: 10.1103/PhysRevLett.111.250404. http://link.aps.org/doi/10.1103/PhysRevLett.111.250404.

[4] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Physical Review Letters*, pages 105–108, 1989. ISSN 00319007. doi: 10.1103/PhysRevLett.63.105. http://link.aps.org/doi/10.1103/PhysRevLett.63.105.

[5] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *Journal of Statistical Physics*, 104(3-4):817–879, 2001. ISSN 00224715. doi: 10.1023/A:1010388907793. http://dx.doi.org/10.1023/A%3A1010388907793.

[6] H. Leff and A. F. Rex. *Maxwell's Demon 2: Entropy, Classical and Quantum Information, Computing*. Taylor & Francis, 2002. ISBN 9780750307598.

[7] C. H. Bennett. The thermodynamics of computationa review. *International Journal of Theoretical Physics*, 21 (12):905–940, December 1982. ISSN 0020-7748. doi: 10.1007/BF02084158. http://www.springerlink.com/content/jn7x3365386phn46/.

[8] R. Landauer. The physical nature of information. *Physics Letters A*, 217(4):188–193, 1996. ISSN 0375-9601. doi: 10.1016/0375-9601(96)00453-7. http://www.sciencedirect.com/science/article/pii/0375960196004537.

[9] J. M. R. Parrondo, J. M. Horowitz, and T. Sagawa. Thermodynamics of information. *Nature Physics*, 11(2):131–139, February 2015. ISSN 1745-2473. doi: 10.1038/nphys3230. http://dx.doi.org/10.1038/nphys3230.

[10] A. N. Kolmogorov. On tables of random numbers. *Sankhya: The Indian Journal of Statistics, Series A*, 25 (4):369–376, 1963. ISSN 0976-836X.

[11] R. Haslinger, K. L. Klinkner, and C. R. Shalizi. The computational structure of spike trains. *Neural computation*, 22(1):121–57, January 2010. ISSN 1530-888X. doi: 10.1162/neco.2009.12-07-678. http://www.mitpressjournals.org/doi/abs/10.1162/neco.2009.12-07-678#.Vay3C_mqqko.

[12] J. B. Park, J. Won Lee, J.-S. Yang, H.-H. Jo, and H.-T. Moon. Complexity analysis of the stock market. *Physica A: Statistical Mechanics and its Applications*, 379(1): 179–187, June 2007. ISSN 03784371. doi: 10.1016/j.physa.2006.12.042. http://www.sciencedirect.com/science/article/pii/S0378437107000271.

[13] C. Lu and R. R. Brooks. P2P hierarchical botnet traffic detection using hidden Markov models. In *Proceedings of the 2012 Workshop on Learning from Authoritative Security Experiment Results - LASER '12*, pages 41–46, New York, New York, USA, July 2012. ACM Press. ISBN 9781450311953. doi: 10.1145/2379616.2379622. http://dl.acm.org/citation.cfm?id=2379616.2379622.

[14] W. Thomson. On the Dynamical Theory of Heat, with numerical results deduced from Mr Joule's equivalent of a Thermal Unit, and M. Regnault's Observations on Steam. *Transactions of the Royal Society of Edinburgh*, XX(II):261–288, 1851. http://www.biodiversitylibrary.org/item/126047#page/291.

[15] K. Wiesner, M. Gu, E. Rieper, and V. Vedral. Information-theoretic lower bound on energy cost of stochastic computation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 468:4058–4066, 2012. ISSN 1364-5021. doi: 10.1098/rspa.2012.0173. dx.doi.org/10.1098/rspa.2012.0173.

[16] S. Still, D. A. Sivak, A. J. Bell, and G. E. Crooks. Thermodynamics of Prediction. *Physical Review Letters*, 109(12):120604, September 2012. ISSN 0031-9007. doi: 10.1103/PhysRevLett.109.120604. http://link.aps.org/doi/10.1103/PhysRevLett.109.120604.

[17] G. E. Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721–2726, September 1999. ISSN 1063-651X. doi: 10.1103/PhysRevE.60.2721. http://link.aps.org/doi/10.1103/PhysRevE.60.2721.

[18] A. Gomez-Marin, J. M. R. Parrondo, and C. Van den Broeck. The footprints of irreversibility. *EPL (Europhysics Letters)*, 82(5):50002, June 2008. ISSN 0295-5075. doi: 10.1209/0295-5075/82/50002. http://stacks.iop.org/0295-5075/82/i=5/a=50002.

[19] M. Gu, K. Wiesner, E. Rieper, and V. Vedral. Quantum mechanics can reduce the complexity of classical models. *Nature Communications*, 3(762), 2012. ISSN 2041-1723. doi: 10.1038/ncomms1761. http://dx.doi.org/10.1038/ncomms1761M3.

[20] J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield. Information accessibility and cryptic processes. *Journal of Physics A: Mathematical and Theoretical*, 42 (36):362002, September 2009. ISSN 1751-8113. doi: 10.1088/1751-8113/42/36/362002. http://stacks.iop.org/1751-8121/42/i=36/a=362002.

APPENDIX

Stationarity, synchronization and determinism. If a process is *stationary*, it follows that $\mathrm{P}(X^{t_1}) = \mathrm{P}(X^{t_2})$ for all t_1, t_2 and hence $H(X^{t_1}) = H(X^{t_2})$ for all t_1, t_2. Any prescient states derived from this sequence then also satisfy $H(R^{t_1}) = H(R^{t_2})$ for all t_1, t_2: the memory required to store the current position in the pattern does not change with time.

If we can *synchronize* our memory, it means that we can determine which prescient state we are in from a sufficiently long observation of the past (at least as many observations as the *crypticity* of the pattern [20]). Expressed in terms of conditional entropy, $H(R^t \,|\, \tilde{X}) = 0$, which reflects no uncertainty in the state R^t if the entire past is known.

Closely related is the concept of *determinism*, which is when the state of the internal memory at the end of any action is completely fixed by the initial state of the memory and the symbols observed in the pattern. Entropically expressed, $H(R^{t+k} \,|\, R^t X^{t+1} \ldots X^{t+k}) = 0$. Once a process has been synchronized to deterministic memory, it will remain synchronized. Causal states are automatically deterministic [5].

Work cost of writing a pattern to a tape. Let us specifically examine the cost of writing k parts of some pattern to a tape using a generator with prescient states R. Initially, all k positions on the tape will be in their default uncorrelated state $X_{\mathrm{dflt}} \ldots X_{\mathrm{dflt}}$ with associated entropy $kH(X_{\mathrm{dflt}})$. For an outside observer who is not aware of the particular choice of symbols written by the machine (but who knew the machine's initial internal state) after the pattern has been written, the entropy of the tape will have changed to $H(X^{t+1} \ldots X^{t+k} \,|\, R^t)$. This value can be expanded to $H(X^{t+1} \,|\, R^t) + H(X^{t+2} \,|\, R^t X^{t+1}) + \ldots + H(X^{t+k} \,|\, R^t X^{t+1} \ldots X^{t+k-1})$.

As R is prescient, R^{t+j} is exactly as useful as $R^t X^{t+1} \ldots X^{t+j}$ for predicting the future of X. Hence, the entropy of the tape can be re-written as $H(X^{t+1} \,|\, R^t) + H(X^{t+2} \,|\, R^{t+1}) + \ldots + H(X^{t+k} \,|\, R^{t+k-1})$. Using stationarity, the entropy of the tape can then be rewritten as $kH(X^{t+1} \,|\, R^t)$. Finally, we note that all prescient states are as good as each other for predicting the future, and so we can replace $H(X^{t+1} \,|\, R^t)$ with the equivalent entropy conditioned on causal states $H(X^{t+1} \,|\, S^t)$. The work cost of writing this pattern is hence proportional to this change in entropy:

$$\beta\, W^k_{\mathrm{tape}} = k \left(H(X_{\mathrm{dflt}}) - H(X^{t+1} \,|\, S^t) \right). \tag{6}$$

Work cost of updating the generator's memory. The above generator must update its internal state from R^t to R^{t+k} before it can produce the next part of the pattern. As with the extractor, we admit an ancilla system of the same size as the internal memory R, which starts and ends in the pure state R_{dflt}. (See fig. 3.) For deterministic memory $H(R^{t+k} \,|\, R^t X^{t+1} \ldots X^{t+k}) = 0$ and the ancilla can be set to state R^{t+k} by a reversible operation with no associated work exchange. For indeterministic memory, the ancilla is slightly randomised and its entropy increases by $H(R^{t+k} \,|\, R^t X^{t+1} \ldots X^{t+k})$ (as with the extractor). Let us consider the cost of resetting the previous internal state R^t back to R_{dflt}. Again, this does not cost the full entropy $H(R^t)$, since there may be some mutual information $I(R^t \,;\, R^{t+k} X^{t+1} \ldots X^{t+k})$ between the old state, and the new state and tape. [Constructively one could imagine this mutual information as being used to reduce the entropy of R^t without work cost, transforming R^t into a new state $\tilde{R}^t$ with entropy $H(R^t \,|\, R^{t+k}, X^{t+1} \ldots X^{t+k})$.] Remarking that, as before, the increased entropy of R^t due to indeterminism is completely negated by the work that could have been extracted by randomizing the ancilla in the first place, we see that the work cost of resetting R^t to R_{dflt} is

$$\beta W^k_{\mathrm{diss}} = H(R^t \,|\, X^{t+1} \ldots X^{t+k} R^{t+k}). \tag{7}$$

We first remark that in the limit of $k \to \infty$, the information $I(R^t \,;\, X^{t+1} \ldots X^{t+k} R^{t+k}) \to E$, the excess entropy, because R is prescient and so contains all the information that could be known about the future. Hence the work dissipated converges on $H(R^t) - E$ in the limit of large k (in agreement with [15] for causal states).

Secondly, by expanding the joint entropy $H(R^t X^{t+1} \ldots X^{t+k} R^{t+k})$ in two different ways, we can rewrite the dissipative entropy as

$$\begin{aligned}
H(R^t \,|\, X^{t+1} \ldots X^{t+k} R^{t+k}) &= H(R^t X^{t+1} \ldots X^{t+k} R^{t+k}) - H(R^{t+k}) - H(X^{t+1} \ldots X^{t+k} \,|\, R^{t+k}) \\
&= H(R^t) + H(X^{t+1} \ldots X^{t+k} \,|\, R^t) + H(R^{t+k} \,|\, R^t X^{t+1} \ldots X^{t+k}) \\
&\quad - H(R^{t+k}) - H(X^{t+1} \ldots X^{t+k} \,|\, R^{t+k}) \\
&= H(X^{t+1} \ldots X^{t+k} \,|\, R^t) - H(X^{t+1} \ldots X^{t+k} \,|\, R^{t+k}) \qquad (8) \\
&= I(X^{t+1} \ldots X^{t+k} \,;\, R^{t+k}) - I(X^{t+1} \ldots X^{t+k} \,;\, R^t). \qquad (9)
\end{aligned}$$

[We have used $H(R^t) = H(R^{t+k})$ from stationarity; and since the term $H(R^{t+k} \,|\, R^t X^{t+1} \ldots X^{t+k})$ arising from indeterminism can be perfectly cancelled out, we have omitted it from the expansion.]

Work extractable from a pattern. Consider an extractor whose internal memory begins in state R^t. For ease of calculation, we admit an ancillary variable of the same size as the internal memory, that is initially in a pure state R_{dflt}. If the device is deterministic, upon the input of k values on the tape, the joint system can be put into the state $R^t X^{t+1} \dots X^{t+k} R^{t+k}$ reversibly and hence without any work cost. If the device is not-deterministic, but still prescient for X, then the final R^{t+k} has some additional randomness, constrainted by the fact R^t is a fine-graining of the causal state. As such, the entropy of the ancilla is increased to $H(R^{t+k} \,|\, S^{t+k})$. We can rewrite this value as $H(R^{t+k} \,|\, R^t X^{t+1} \dots X^{t+k})$, since the causal state is itself determined by $R^t X^{t+1} \dots X^{t+k}$, and no other parts of the past are useful for determining R^{t+k}.

We can now evaluate the entropy of the previous internal state and the tape, noting its mutual information with the new internal state: $I(R^t X^{t+1} \dots X^{t+k} ; R^{t+k}) = H(R^{t+k}) - H(R^{t+k} \,|\, R^t X^{t+1} \dots X^{t+k})$.

$$\begin{aligned} H(R^t X^{t+1} \dots X^{t+k} \,|\, R^{t+k}) & \\ &= H(R^t X^{t+1} \dots X^{t+k}) - I(R^t X^{t+1} \dots X^{t+k} ; R^{t+k}) \\ &= H(R^t) + H(X^{t+1} \,|\, R^t) + \dots \\ &\quad + H(X^{t+k} \,|\, R^t X^{t+1} \dots X^{t+k-1}) \\ &\quad - H(R^{t+k}) + H(R^{t+k} \,|\, R^t X^{t+1} \dots X^{t+k}) \\ &= kH(X^{t+1} \,|\, S^t) + H(R^{t+k} \,|\, R^t X^{t+1} \dots X^{t+k}) \end{aligned} \tag{10}$$

We have used that for a stationary process $H(R^{t+k}) = H(R^t)$. We have also used prescience, such that conditioning the future of X on $R^t X^{t+1} \dots X^{t+j}$ is equivalent to conditioning it on the causal state S^{t+j}. Since the terms in S^{t+j} are shifted in time with respect to each other, from stationarity they all represent the same value, and so each addend contributes the same entropy.

The extractor takes each value on the tape back to the default uncorrelated states X_{dflt}, with entropy $kH(X_{\text{dflt}})$. The final state of the ancilla must also be pure and uncorrelated with the reset tape, such that the total final entropy of the tape-ancilla system is also $kH(X_{\text{dflt}})$. This corresponds to a change in entropy $k\left[H(X_{\text{dflt}} - H(X^{t+1} \,|\, S^t))\right] - H(R^{t+k} \,|\, R^t X^{t+1} \dots X^{t+k})$. The final term, arising only when there is indeterminism, is equal in magnitude but opposite in sign to the amount of randomness we introduced to the tape from indeterminism in the first place. As such, these terms perfectly cancel out, and do not affect the final value of the work exchange. By Landauer's principal, we hence see that the work extracted when resetting k symbols on the tape and advancing the internal state from R^t to R^{t+k} is

$$\beta W^k_{\text{out}} = k\left[H(X^{t+1} \,|\, S^t) - H(X_{\text{dflt}})\right]. \tag{11}$$